Staff Analytical Note/Note analytique du personnel — 2021-21

Last updated: September 9, 2021

Canadian housing supply elasticities

by Nuno Paixão

Financial Stability Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9

nmarquesdapaixao@bankofcanada.ca

ISSN 2369-9639

Introduction

Housing price dynamics have been a source of intense debate and concern in several economies, including Canada's. What drives growth in house prices? Why do house prices evolve differently across cities? To answer these questions, we need to consider both demand and supply factors. On the demand side, we can highlight, among other things, household income, financing costs and population change. On the supply side, construction costs and local housing supply elasticities are key elements in explaining the evolution of house prices.

This note focuses on housing supply elasticities. A city with a more inelastic housing supply may face higher increases in house prices in reaction to positive local or aggregate shocks. Therefore, estimates of local housing supply elasticities are crucial for predicting changes in house prices and evaluating the impact of multiple policy interventions.

This note delivers the first estimates of the housing supply elasticities for the largest cities in Canada (census agglomerations, CA[1]) following the approach developed by Guren et al. (2021).[2] This novel approach exploits the fact that house prices in some cities are systematically more sensitive to regional cycles than house prices in other cities. This approach differs from the one used by Saiz (2010) in estimating housing supply elasticities for most metropolitan areas in the United States. He does this by exploiting city-specific building regulations and land unavailability, specifically the land within a 50-kilometre radius of the city centre that is not suitable for construction due to geographic constraints such as steep slopes or bodies of water. His estimates are widely used in the economic literature in model calibrations and as an instrument for the change in house prices during the boom and bust cycle of the 2000s.

The approach developed by Guren et al. (2021) has two main advantages over the one of Saiz (2010). First, the Saiz measure correlates with other city characteristics such as productivity and growth in demand (Davidoff, 2016). This raises the concern that higher house price volatility in some cities is not driven by inelastic housing supply, as estimated by Saiz, but by differences in other characteristics such as different industrial composition and different exposure to secular trends, for example, an increase in housing demand in coastal areas with inelastic supply. To address this shortcoming, Guren et al. (2021) employ a panel specification that allows them to control for city-specific trends, different sensitivity to regional business cycles, and changes in the city's population and industry structure. Second, by exploiting the systematic historical sensitivity of local house prices to regional house price cycles, this new approach allows us to estimate housing supply elasticities without resorting to

[1] https://www150.statcan.gc.ca/n1/pub/92-195-x/2011001/geo/cma-rmr/cma-rmr-eng.htm

[2] Housing supply elasticities for Canada's largest cities were first estimated in Giannone et al. (2020). This note presents these estimates and the procedure used.

geographical and regulation data across Canadian cities, data that are not currently available for most cities in Canada.

Chart 1: House prices in Vancouver, Winnipeg and the West region

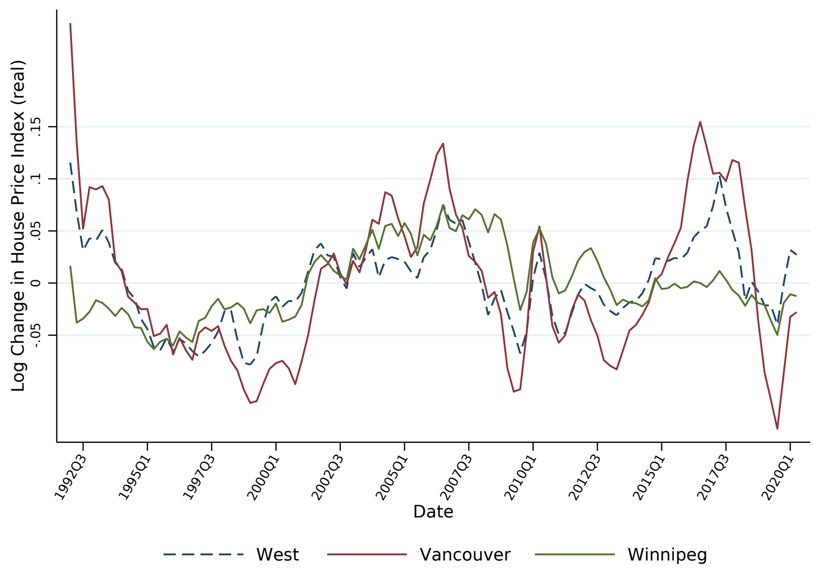

Note: All time series correspond to the annual log change in the House Price Index. All series are demeaned relative to the city or region average. The West region includes all provinces west of Ontario. Source: Teranet

Guren et al. (2021) estimate housing supply elasticities by exploiting systematic differences in cities' responses to regional house price cycles. Sinai (2012) documents that house prices in some US cities are systematically more sensitive to regional cycles than those in other cities, which is also true for Canada. Let's consider Vancouver and Winnipeg. Chart 1 plots the annual log change of real house prices in the West,[3] Vancouver and Winnipeg from 1992 to 2020. The West region experienced several regional boom-bust cycles throughout the sample period. Vancouver and Winnipeg also experienced several cycles that tended to correlate with the regional ones. However, house prices in Vancouver tended to increase by more than house prices in Winnipeg when regional house prices were booming. They also decreased by more when regional prices were contracting.

This systematic difference in the sensitivity of house prices in different cities to the regional house price cycles is crucial for the identification strategy described in the next section.

[3]The West region includes all provinces west of Ontario.

Empirical strategy

A simple approach to estimating the sensitivity of house prices in different cities to regional house price movements, γ_i, consists of running the regression:

$$\Delta p_{i,r,t} = \phi_i + \chi_{r,t} + \gamma_i \Delta P_{r,t} + \epsilon_{i,r,t} \tag{1}$$

where $\Delta p_{i,r,t}$ denotes the log annual change of real house prices of city i in region r, and $\Delta P_{r,t}$ stands for the log annual change in regional house prices.[4] This specification includes city fixed effects, ϕ_i, to control for unobserved city heterogeneity, and region-time fixed effects, $\chi_{r,t}$, to control for trends at the regional level. Cities with higher $\hat{\gamma}_i$, the estimate of γ_i, are cities that systematically respond to regional shocks with higher fluctuations in higher prices and, therefore, cities with more inelastic housing supply. Therefore, $\hat{\gamma}_i$ denotes the proxy for the inverse of the housing supply elasticity.

This simple approach, however, assumes that local house prices respond differently to regional house price shocks only because of differences in the housing supply elasticity. This assumption seems too restrictive because differences in the structure of the local economy may cause different responses. Applying the example in Guren et al. (2021) to the Canadian context, we suppose that Vancouver has an industrial structure tilted toward highly cyclical durable goods relative to that of Winnipeg. A positive aggregate demand shock would consequently lead to higher increases in employment and house prices in Vancouver than in Winnipeg. Therefore, γ_i would be estimated to be higher in Vancouver than in Winnipeg purely due to reverse causality. Then, variation in $\hat{\gamma}_i$ would reflect not only differences in housing supply elasticities across cities but also potentially other confounding factors.

To address these concerns, I apply a refined version of equation (1) similar to the one proposed by Guren et al. (2021):

$$\Delta p_{i,r,t} = \phi_i + \gamma_i \Delta P_{i,r,t} + \delta_i \Delta y_{i,r,t} + \mu_i \Delta Y_{r,t} + \Gamma X_{i,r,t} + \epsilon_{i,r,t} \tag{2}$$

This version augments equation (1) with local and regional changes in per capita retail, construction and manufacturing employment with city-specific coefficients. The vectors with these changes in employment at the city and regional levels are $\Delta y_{i,r,t}$ and $\Delta Y_{r,t}$, respectively. This specification controls for the different impact across cities of different demand shocks reflected in these industries.[5] It also includes another set of controls, $X_{i,r,t}$, specifically

[4]Throughout this note, I follow the same notation simplification as in Guren et al. (2021), where $\gamma_i \Delta P_{r,t}$ is used to denote $\sum_i \gamma_i \Delta P_{i,r,t} I_i$, where I_i is an indicator for city i.

[5]Guren et al. (2021) control for local and regional changes in retail employment only. The correlation between the baseline estimates and the γ_i estimates using this less strict specification is 97 percent. If instead

two-digit industry code shares multiplied by time dummies. This structure allows for non-parametrically controlling for all variation that is correlated with industry structure in the cross-section. I also depart from Guren et al. (2021) by controlling for population growth at the city and regional levels and for real mortgage rates.

Overall, this refined approach implies that $\hat{\gamma}_i$ is estimated using local house price variation that is independent of local and regional changes in employment and all other controls included in $X_{i,r,t}$. It is therefore not subject to the bias resulting from the reverse causality explained before. The key identifying assumption is that, conditional on controls, there are no other aggregate factors sensitive to house prices as captured by γ_i that are correlated with regional house prices in the time series and that differentially impact employment per capita in the same city. However, this approach does not require exogenous variation in regional house prices. Common factors can drive regional house prices, regional economic activity and even local prices and activity.

Data

I estimate the elasticities using the House Price Index developed by Teranet at the forward sortation area (FSA) level. City and regional house prices are built by aggregating them using the 2011 FSA populations as weights.[6] House prices are converted into a real index by using the gross domestic product (GDP) deflator from Statistics Canada. Following the methodology in Guren et al. (2021), I consider quarterly house prices and calculate annual changes of the log of the House Price Index in real terms. Annual two-digit industry employment codes and population at the CA level are obtained from Statistics Canada. Real monthly mortgage rates are from the Bank of Canada.

Guren et al. (2021) consider four regions when estimating the elasticities for US cities. Given the significant differences in size between Canada and the United States, I consider three regions in Canada: east, west and northern territories.

of controlling for changes in employment per capita in these three industries separately, I control only for aggregate changes in per capita employment, the correlation drops to 95 percent. More importantly, if I don't control for any changes in the employment growth across different industries at the city level, the correlation drops to 23 percent, which reflects the importance of controlling for changes in industry composition.

[6] Very similar results are obtained if total dwellings or total occupied dwellings are used as weights.

Results

Chart 2 plots the estimated housing supply elasticities for Canadian census metropolitan areas (CMAs), specifically, the inverse of $\hat{\gamma}_i$ estimated from equation (2).[7] The median housing supply elasticity is 2.2 among all CAs and 1.94 if I restrict the sample to CMAs. These estimates imply that a 1 percent increase in house prices in the median Canadian city is associated with an increase in housing supply of 2.2 percent. Alternatively, we can think that, all else equal, a 1 percent increase in housing demand leads to an increase in house prices in the median city of 0.45 percent (1/2.2).

Chart 2: Housing supply elasticities for Canadian census metropolitan areas

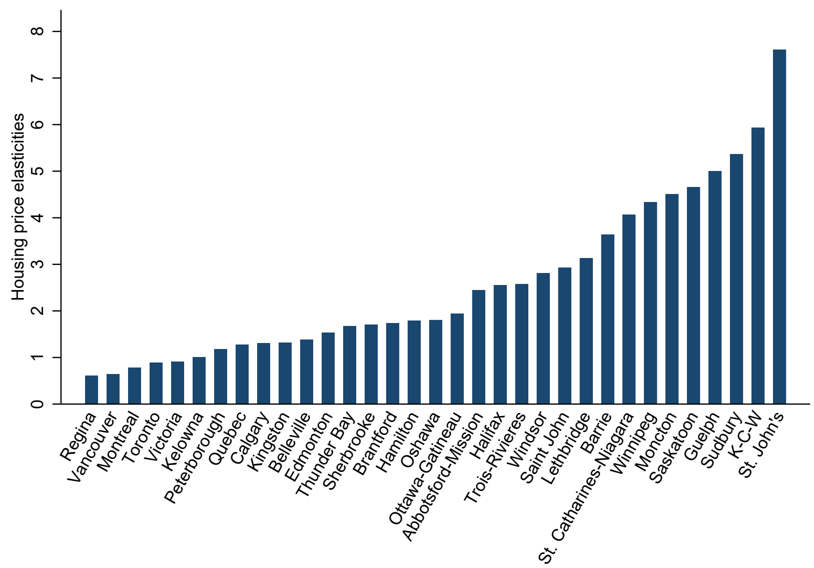

Note: This chart plots the estimated housing supply elasticity, $1/\hat{\gamma}_i$, estimated from equation (2). London and Saguenay are excluded for vizualization purposes. The elasticities for these two CMAs are 19.6 and 21.6, respectively. K-C-W stands for Kitchener-Cambridge-Waterloo.

Chart 2 also shows significant heterogeneity across cities. Going back to the previous example, elasticities in Vancouver and Winnipeg are 0.63 and 4.34, respectively. Assuming that both cities face a 1 percent increase in housing demand, house prices are predicted to increase 1.57 percent in Vancouver and 0.23 percent in Winnipeg. For comparison, the median housing supply elasticity among US metropolitan areas estimated in Saiz (2010) is 2.26, very close to the median elasticity in Canada. Saiz also estimates elasticities of 0.63 and 0.72 in New York and San Francisco, respectively. These values compare closely with Vancouver and

[7]The procedure estimates housing supply elasticities for 151 CAs, but for clarity the chart is restricted to the CMAs. For better visualization, London and Saguenay are also excluded from the chart. The elasticities for these two CMAs are 19.6 and 21.6, respectively.

Toronto, where estimated elasticities in this note are 0.64 and 0.89, respectively. However, the distribution of elasticities in Canada is more skewed to the right than it is in the United States. A larger share of cities in Canada have very elastic housing supplies.

Housing supply elasticities and long-run growth in house prices

As stated earlier, the dynamics in the housing market depend on both demand and supply. Despite the importance of housing supply elasticities, they alone cannot explain dynamics in this market, especially in the short run. Nevertheless, we can see that housing supply elasticities correlate with long-run growth in house prices.[8] Chart 3 shows the weighted average of the Teranet House Price Index across cities in Canada between 1990 and 2020.[9] The sample is split between cities with housing supply below and above the median.

Chart 3: Housing supply elasticities and house prices

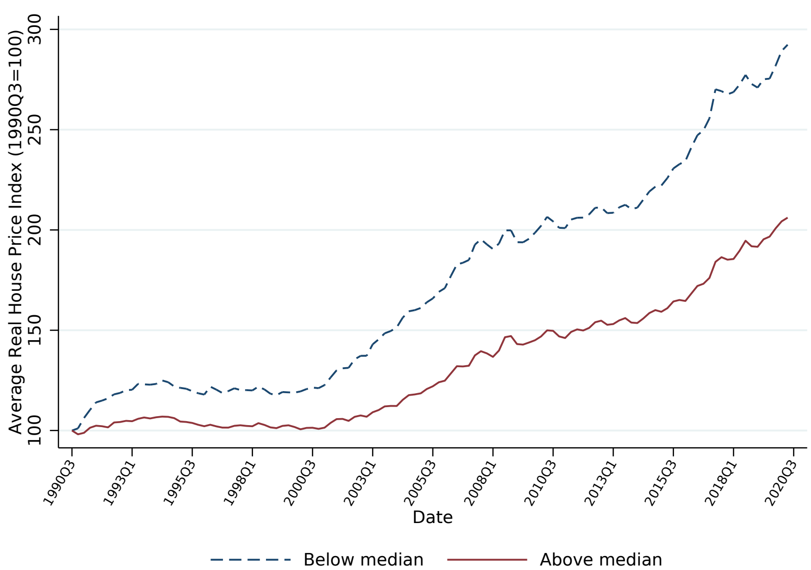

Note: Series are weighted by population at the city level from Statistics Canada. The House Price Index has been normalized by the gross domestic product deflator. House price indexes are normalized to 100 in 1990Q3 to highlight the distinct growth rates between the two groups of cities. Source: Teranet

We can see that cities in Canada with more inelastic housing supply (elasticity below median) faced a higher house price growth during this period than the cities with more elastic

[8]To identify the importance of housing supply elasticities in explaining the different long-run housing market dynamics across cities, we need a structural model where local and aggregate demand shocks are properly identified. This exercise goes beyond the purpose of this note.

[9]Both series are normalized to 100 in the third quarter of 1990 to highlight the distinct growth rates between the two groups of cities.

supply. The average house price index for cities with inelastic supply in 2020 was four times that in 1990. For the cities with more elastic supply, it was only three times that in 1990. This suggests that cities with inelastic housing supply face larger constraints, which can explain higher house price growth.

Conclusion

This note presents estimates of housing supply elasticities for the largest Canadian cities and describes the procedure followed. In contrast with the popular elasticities estimated by Saiz (2010), who explores geographical and regulatory heterogeneity across cities, the approach used in this note follows Guren et al. (2021), who exploit the systematic different sensitivity of different cities to regional house price cycles. This procedure implicitly accounts for diverse factors that determine housing supply elasticities in addition to the geographical and regulatory factors. In particular, it incorporates expectations about long-run growth. Many cities that do not face land shortages currently may become constrained if they experience land shortages in the future.

The median housing supply elasticity among all Canadian CAs is 2.2, but Canada presents a high degree of heterogeneity across cities.

References

Davidoff, T. 2016. “Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated with Many Demand Factors.” *Critical Finance Review* 5 (2): 177–206.

Giannone, E., Q. Li, N. Paixao and X. Pang. 2020. “Unpacking Moving.”

Guren, A. M., A. McKay, E. Nakamura and J. Steinsson. 2021. “Housing Wealth Effects: The Long View.” *Review of Economic Studies* 88 (2): 669–707.

Saiz, A. 2010. “The Geographic Determinants of Housing Supply.” *Quarterly Journal of Economics* 125 (3): 1253–1296.

Sinai, T. M. 2012. “House Price Moments in Boom-Bust Cycles.” National Bureau of Economic Research Working Paper No. 18059.